Exploring Multicore and SDN for high-speed data movement

Wenji Wu, Phil DeMar
Fermilab
May 12, 2014, Washington DC

SDN enables innovative network use cases

- Consistent routing and security policy enforcement [Ethane, SIGCOMM 2007]
- New data center network architecture like VL2 [Tavakoli et al. Hotnets 2009]
- Flow scheduling [Hedera, NSDI 2010]
- Energy efficient data center [ElasticTree, NSDI 2010]
- Automated data center QoS (Kim et al. INM/WREN 2010)
- Traffic engineering for big data applications (Wang et al. HOTSDN 2012)

A NEW SDN USE CASE

We use SDN to improve network I/Os on multicore systems

Data Transfer Node (DTN)

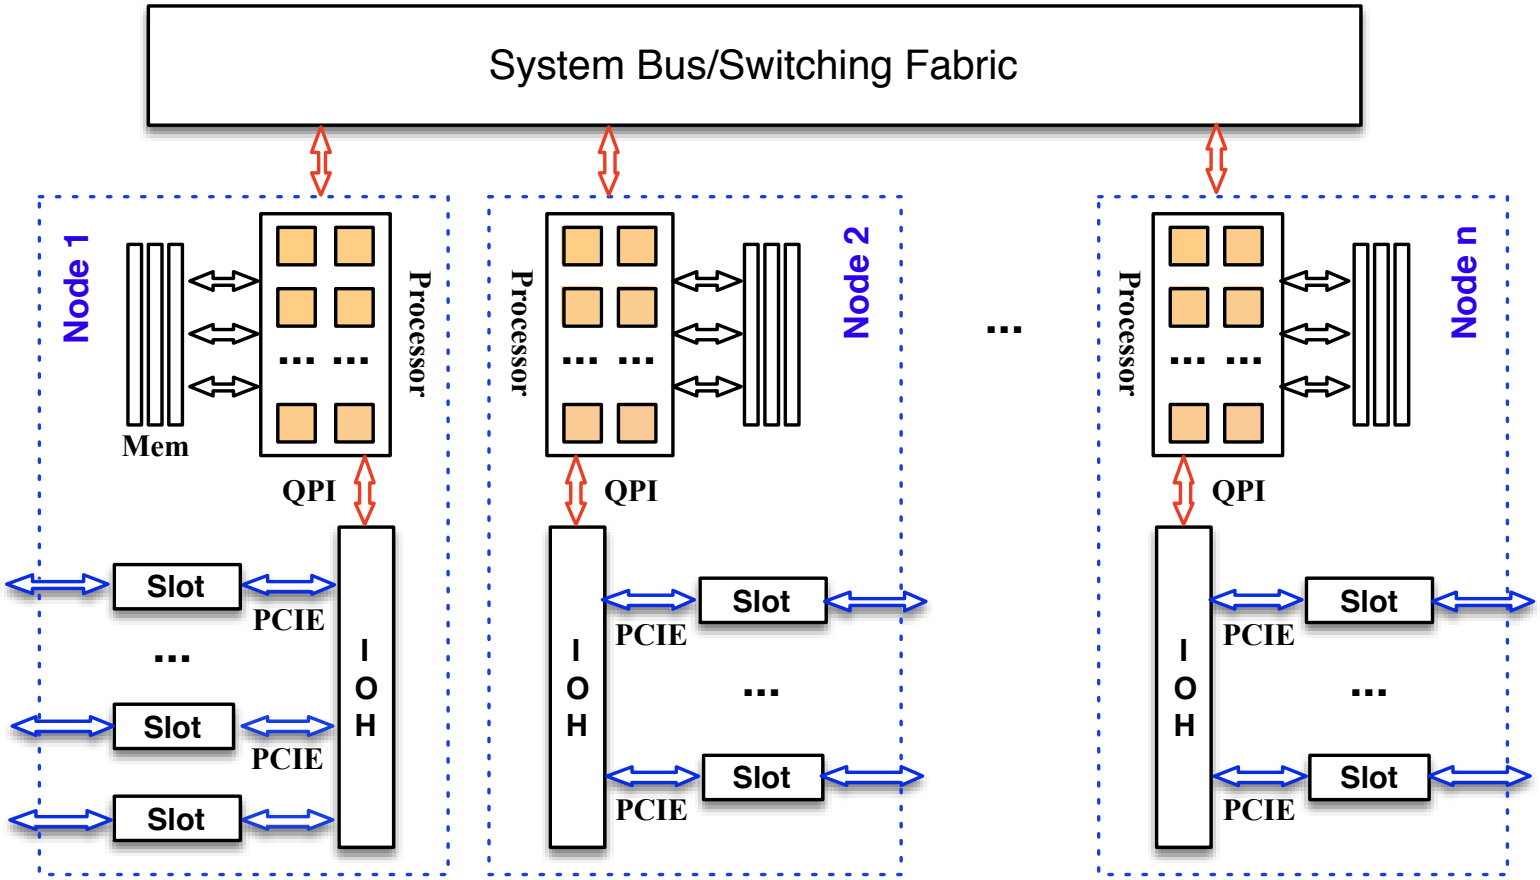

An NUMA system, with multiple NUMA nodes
Each node has one or multiple I/O devices

Principles to achieve high-speed data movement on DTNs

- Maximizing parallelism
- NUMA-awareness
- I/O locality
- Core-affinity on network I/O
- Core-affinity on disk I/O

Which DTN Configuration is better?

DTN configuration 1

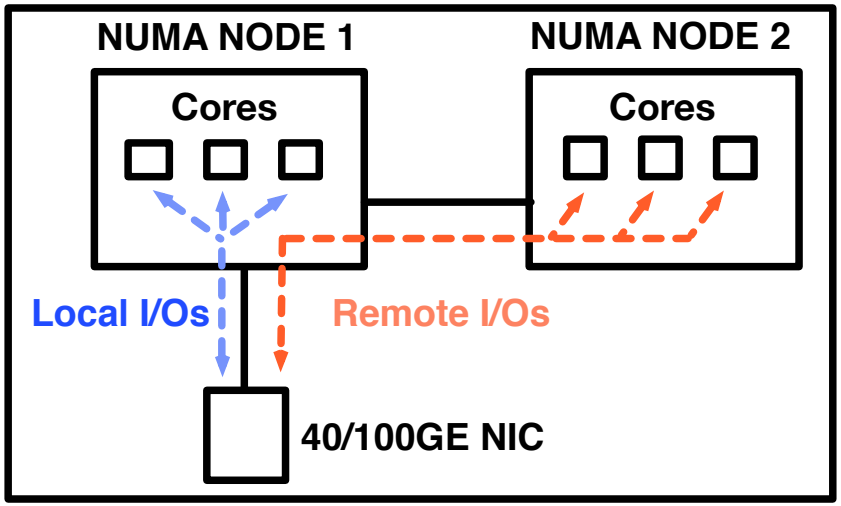

a. With on big NIC (40/100 GE NIC)

DTN configuration 2

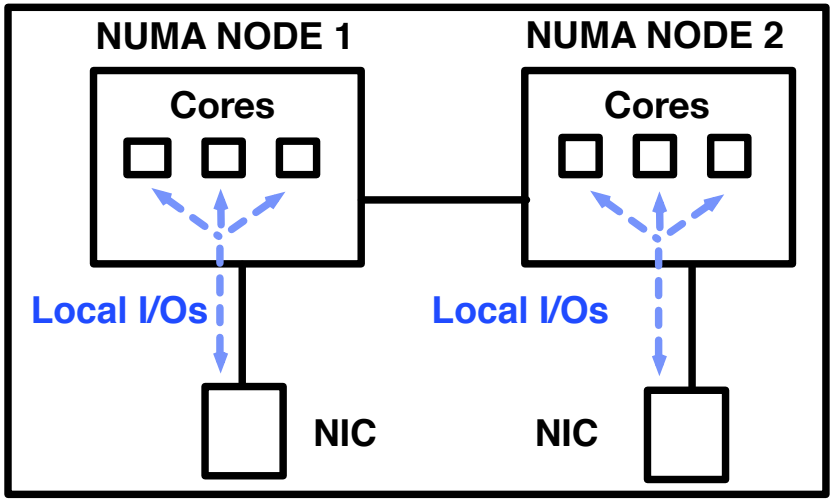

b. With multiple smaller NICs (e.g., 10GE)

Typically one IP address

How to ensure core affinity on network I/Os?

A traditional switch cannot ensure core affinity on network I/O

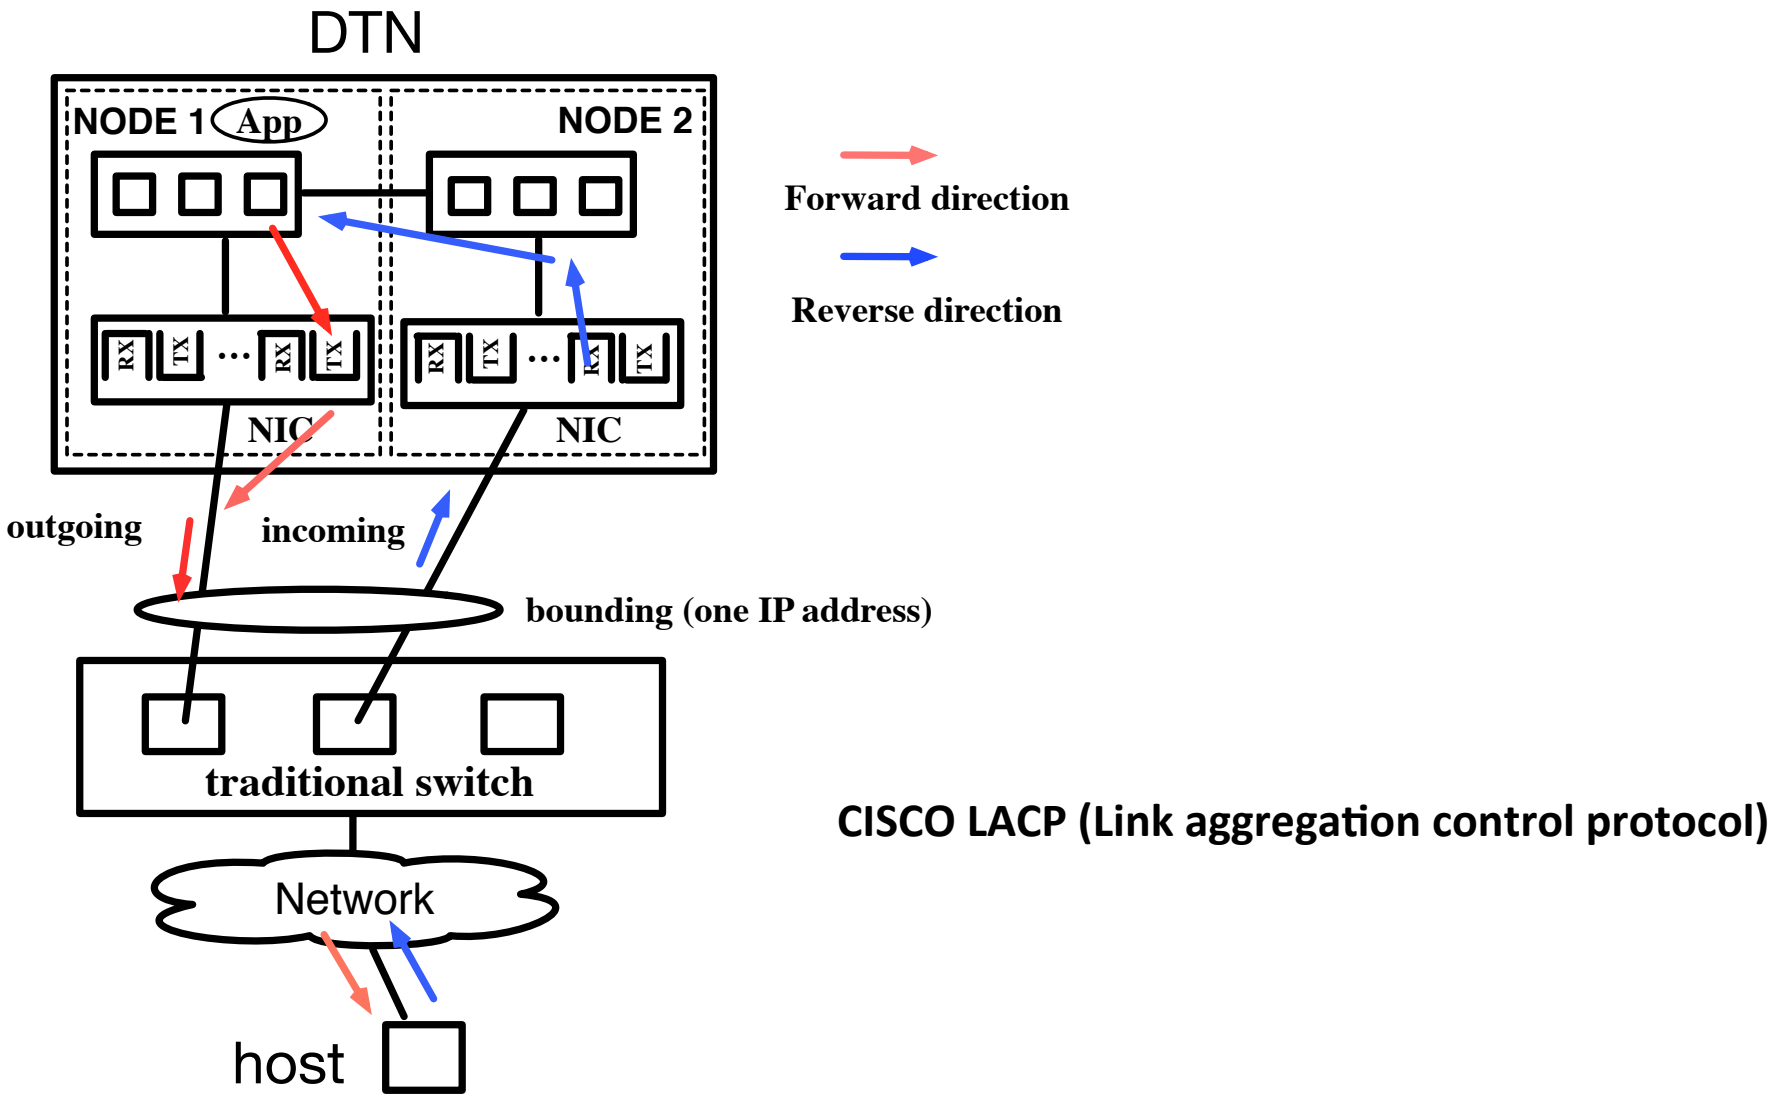

CISCO LACP (Link aggregation control protocol)

One Solution: Use SDN to improve network I/Os on multicore systems

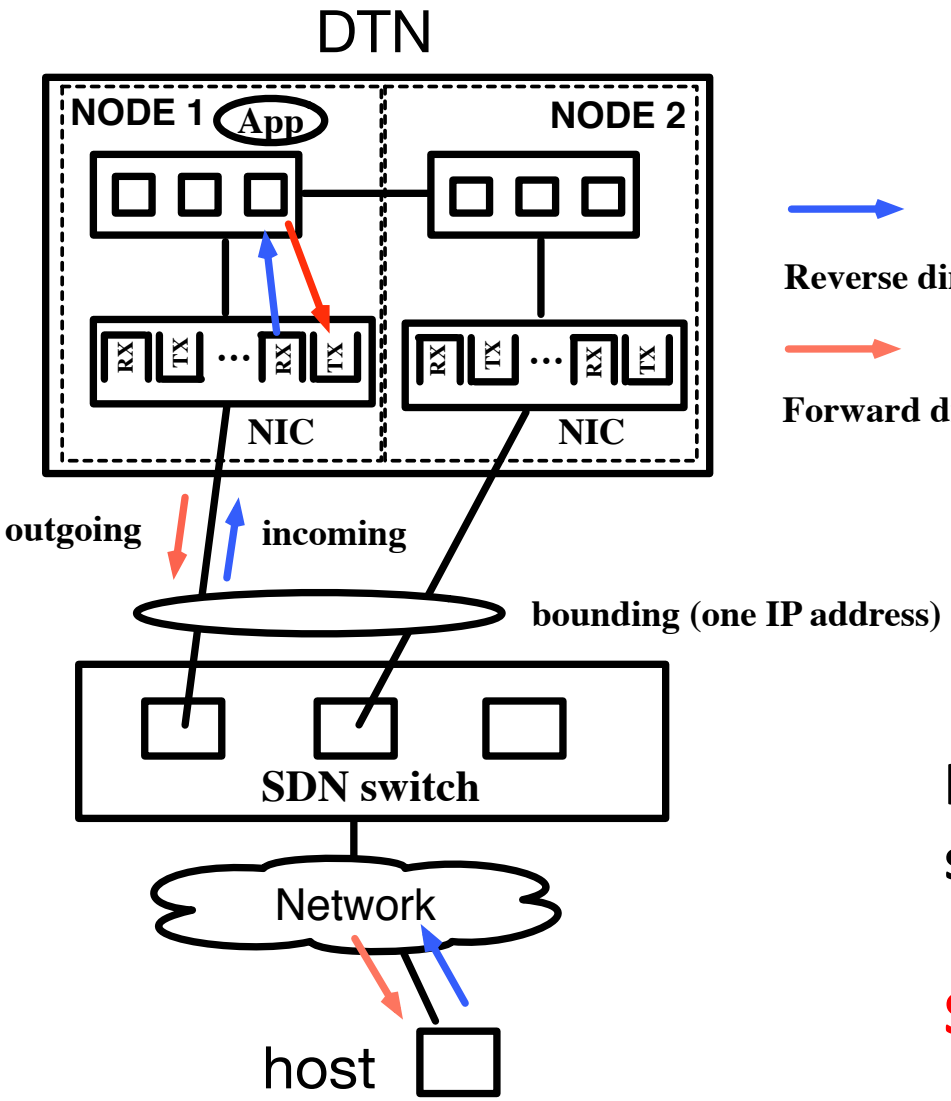

In an SDN network, traffic can be steered on a per-flow basis

SDN used to ensure core affinity